

CLAIMS

What is claimed is:

- 1 1. An apparatus for managing flow control of a data transfer, comprising:
2 a first protocol associated with a plurality of receive buffers;
3 a second protocol adapted to manage the plurality of receive buffers for the first
4 protocol; and
5 a third protocol that determines whether one of the plurality of receive buffers is
6 available for a data packet and (a) if one of the plurality of receive buffers is available,
7 permits an acknowledgement packet to be sent to a node that sent the data packet, and
8 (b) if one of the plurality of receive buffers is unavailable, drops the data packet, notifies
9 the second protocol regarding the unavailability of the plurality of receive buffers, and
10 withholds the acknowledgement packet.

- 1 2. The apparatus set forth in claim 1, wherein the first protocol is an upper
2 layer protocol ("ULP").

- 1 3. The apparatus set forth in claim 2, wherein the upper layer protocol is an
2 internet small computer systems interface ("iSCSI") protocol.

1 4. The apparatus set forth in claim 1, wherein the second protocol is a
2 datamover protocol.

1 5. The apparatus set forth in claim 1, wherein the third protocol is an
2 iWARP protocol.

1 6. The apparatus set forth in claim 5, wherein the iWARP protocol is a
2 direct data placement (“DDP”) protocol.

1 7. The apparatus set forth in claim 1, comprising a transport protocol that
2 generates a request to the third protocol to determine whether one of the plurality of
3 receive buffers is available for the data packet.

1 8. The apparatus set forth in claim 1, wherein the data packet comprises a
2 sequence field that corresponds to a reliability tracking value for the data packet.

1 9. The apparatus set forth in claim 1, wherein the acknowledgement packet
2 comprises an acknowledgement field that corresponds to an identity of data received by
3 the transport protocol.

1 10. The apparatus set forth in claim 1, comprising a transport protocol that
2 uses a remote direct memory access network interface card ("RNIC") to receive the data
3 packet and send the acknowledgement packet.

1 11. A network, comprising:
2 a plurality of systems, at least one of the plurality of systems executing a
3 process; and
4 at least one input/output device adapted to receive a data packet from the at least
5 one of the plurality of systems, the at least one input/output device comprising:
6 a first protocol associated with a plurality of receive buffers;
7 a second protocol adapted to manage the plurality of receive buffers for
8 the first protocol; and
9 a third protocol that determines whether one of the plurality of receive
10 buffers is available for a data packet and (a) if one of the plurality
11 of receive buffers is available, permits an acknowledgement

12 packet to be sent to a node that sent the data packet, and (b) if one
13 of the plurality of receive buffers is unavailable, drops the data
14 packet, notifies the second protocol regarding the unavailability
15 of the plurality of receive buffers, and withholds the
16 acknowledgement packet.

1 12. The apparatus set forth in claim 11, wherein the first protocol is an upper
2 layer protocol ("ULP").

1 13. The apparatus set forth in claim 12, wherein the upper layer protocol is
2 an internet small computer systems interface ("iSCSI") protocol.

1 14. The apparatus set forth in claim 11, wherein the second protocol is a
2 datamover protocol.

1 15. The apparatus set forth in claim 11, wherein the third protocol is an
2 iWARP protocol.

1 16. The apparatus set forth in claim 15, wherein the iWARP protocol is a
2 direct data placement (“DDP”) protocol.

1 17. The apparatus set forth in claim 11, comprising a transport protocol that
2 generates a request to the third protocol to determine whether one of the plurality of
3 receive buffers is available for the data packet.

1
2

3 18. The apparatus set forth in claim 11, wherein the data packet comprises a
4 sequence field that corresponds to a reliability tracking value for the data packet.

1 19. The apparatus set forth in claim 11, wherein the acknowledgement packet
2 comprises an acknowledgement field that corresponds to an identity of data received by
3 the transport protocol.

1 20. The apparatus set forth in claim 11, comprising a transport protocol that
2 uses a remote direct memory access network interface card (“RNIC”) to receive the data
3 packet and send the acknowledgement packet.

1 21. A method of managing flow control of a data transfer, the method
2 comprising the acts of:
3 receiving a data packet;
4 determining whether at least one receive buffer is available for the data packet;
5 if the at least one buffer is available, sending an acknowledgement packet to a
6 node that sent the data packet; and
7 if the at least one buffer is unavailable, dropping the data packet, providing a
8 notification regarding the unavailability of the at least one buffer, and
9 withholding an acknowledgement packet from the node that sent the data
10 packet.

1 22. The method set forth in claim 21, comprising the act of placing the data
2 packet into the at least one buffer if the at least one buffer is available.

1 23. The method set forth in claim 21, comprising the act of transmitting the
2 data packet according to a transmission control protocol ("TCP").

1 24. The method set forth in claim 21, comprising the act of providing the
2 notification regarding the unavailability of the at least one buffer via an internet small
3 computer systems interface ("iSCSI") protocol.

1 25. The method set forth in claim 21, comprising the act of notifying a
2 process associated with the at least one buffer once the at least one buffer is determined
3 to be unavailable.

1 26. An apparatus for managing flow control of a data transfer, comprising:
2 means for receiving a data packet at a first protocol;
3 means for determining whether at least one receive buffer is available for the
4 data packet;
5 means for sending an acknowledgement packet to a node that send the data
6 packet if the at least one buffer is available; and
7 means for dropping the data packet, notifying a second protocol regarding the
8 unavailability of the at least one buffer, and preventing an acknowledgement packet
9 from being sent if the at least one buffer is unavailable.